

**METHOD AND SYSTEM FOR ABSORBING
DEFECTS IN HIGH PERFORMANCE MICROPROCESSOR
WITH A LARGE N-WAY SET ASSOCIATIVE CACHE**

5

CROSS-REFERENCE TO RELATED APPLICATIONS

This application relates to the following commonly assigned co-pending applications entitled:

- 10 “Scan Wheel – An Apparatus For Interfacing A High Speed Scan-Path With A Slow Speed Tester,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-23700; “Rotary Rule And Coherence Dependence Priority Rule,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27300; “Speculative Scalable Directory Based Cache Coherence Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27400;
- 15 “Scalable Efficient IO Port Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27500; “Efficient Translation Buffer Miss Processing For Applications Using Large Pages In Systems With A Large Range Of Page Sizes By Eliminating Page Table Level,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27600; “Fault Containment And Error Recovery Techniques In A Scalable Multiprocessor,” Serial No.
- 20 _____, filed August 31, 2000, Attorney Docket No. 1662-27700; “Speculative Directory Writes In A Directory Based CC-Non Uniform Memory Access Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27800; “Special Encoding Of Known Bad Data,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27900; “Broadcast Invalidate Scheme,” Serial No. _____, filed August 31, 2000, Attorney Docket
- 25 No. 1662-28000; “Mechanism To Keep All Pages Open In A DRAM Memory System,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-28100; “Programmable

DRAM Address Mapping Mechanism,” Serial No. _____, filed August 31, 2000, Attorney
Docket No. 1662-28200; “Mechanism To Enforce Memory Read/Write Fairness, Avoid Tristate
Bus Conflicts, And Maximize Memory Bandwidth,” Serial No. _____, filed August 31,
2000, Attorney Docket No. 1662-29200; “An Efficient Address Interleaving With Simultaneous
5 Multiple Locality Options,” Serial No. _____, filed August 31, 2000, Attorney Docket No.
1662-29300; Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-29400; “A
Method For Reducing Directory Writes And Latency In A High Performance, Directory-Based,
Coherency Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-
29600; “Mechanism To Reorder Memory Read And Write Transactions For Reduced Latency And
10 Increased Bandwidth,” Serial No. _____, filed August 31, 2000, Attorney Docket
No. 1662-30800; “Look-Ahead Mechanism To Minimize And Manage Bank Conflicts In A
Computer Memory System,” Serial No. _____, filed August 31, 2000, Attorney Docket
No. 1662-30900; “Resource Allocation Scheme That Ensures Forward Progress, Maximizes
Utilization Of Available Buffers And Guarantees Minimum Request Rate,” Serial
15 No. _____, filed August 31, 2000, Attorney Docket No. 1662-31000; “Input Data Recovery
Scheme,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-31100; “Fast
Lane Prefetching,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-
31200; “Mechanism For Synchronizing Multiple Skewed Source-Synchronous Data Channels
With Automatic Initialization Feature,” Serial No. _____, filed August 31, 2000, Attorney
20 Docket No. 1662-31300; “A Mechanism To Control The Allocation Of An N-Source Shared
Buffer,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-31400; and
“Chaining Directory Reads And Writes To Reduce DRAM Bandwidth In A Directory Based CC-

NUMA Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-31500, all of which are incorporated by reference herein.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

Not applicable.

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to a method and architecture for absorbing defects and improving the yield of a microprocessor having a large on-chip cache. More particularly, the invention relates to improving the yield of a microprocessor having a large on-chip n-way set associative cache by absorbing or working around defects in the portion of the die allocated to cache.

Background of the Invention

In general, when designing microprocessor-based systems, system performance can be enhanced by increasing the random access memory ("RAM") cache available on-chip to the microprocessor. This is because accessing on-chip cache is significantly faster than accessing other off-chip memory, such as single inline memory modules ("SIMMs") or dual inline memory modules ("DIMMs"). So, at the risk of over-simplifying, the more on-chip cache available the better.

The problem is that increasing available on-chip cache results in increasing the die size for the microprocessor. As the size of the die increases, generally the manufacturing yields for the die

decrease. In fact, typically the yield goes down exponentially as the die size is increased. This means that it is harder to manufacture large dies that are not defective.

This creates two competing interests in the design of microprocessors. On the one hand, one would like as much cache as possible available on-chip to increase the speed and efficiency of the microprocessor. On the other hand, any increase in the die size will probably result in reduced production yields for the microprocessor. Industry testing has indicated that for up to about 4 megabytes of cache, the return on speed and efficiency is often worth the resultant manufacturing issues. After that cache size, however, there may be diminishing returns. That is, the benefits of the increased cache size may be outweighed by the reduction in manufacturing yields. Ultimately, a general rule would be that one wants as much cache as can fit on the die while maintaining acceptable production yields.

On typical microprocessor dies, then, large areas of the die are allocated to the cache. In fact, the cache typically takes up more physical real estate on the die than anything else. This necessarily means that manufacturing defects in a given microprocessor will often occur in the cache portion of the die since it is the largest physical portion of the die. Accordingly, if there was some way to organize and manage the cache to work around these defects, production yields could be increased. Any method or system that increases the number of defects which a die can absorb while still functioning properly will have a significant yield benefit.

The state of the art currently provides for segmenting the data array of the cache to allow the cache to absorb or "work around" some defects in the data array of the cache. In particular, segmenting the data array of the cache allows for some redundancy and selectivity in the data array that allows the cache to work around some unrepairable defects. For example, by assigning rows and columns to the data array of the cache, row and column redundancy can be used to replace

defective rows or columns of the data array. That is, where a particular row or column is found to have an unrepairable defect, it can be replaced with one of the redundant rows or columns that is not defective. Additionally, in a set associative cache where the data array is divided into a plurality of sets or ways, any way found to have a defect can be disabled. This allows an otherwise defective die to still be used, although with a smaller usable cache.

The present invention is directed at a method and architecture for working around defects in a set associative cache, thereby allowing larger on-chip cache while maintaining acceptable manufacturing yields. The present invention can be used in combination with other methods, such as row and column redundancy, to further increase yields.

BRIEF SUMMARY OF THE INVENTION

In accordance with the present invention, there is provided a novel method and architecture for increasing the number of defects in the data array of the cache which can be absorbed while maintaining a useable cache size thereby reducing the percentage of dies which must be discarded due to manufacturing defects. This is accomplished by remapping defective portions of ways in a set associative cache to a surrogate portion of another way in the cache. By utilizing a multiplexer or comparable switching mechanism ("mux") in the shortest path between the access control logic of the microprocessor and the closest way, additional selectivity can be gained. More specifically, the mux allows smaller portions of the way to be disabled and replaced with a useable portion of a surrogate way, *i.e.*, the way with the shortest path. Since the surrogate way has the shortest physical path, the mux can be added without adding any latency or cycle time. This allows for a larger percentage of die to be repaired, with larger useable cache remaining.

The inventive architecture for set associative cache comprises: a set associative cache having a plurality of ways wherein the ways are segmented into a plurality of banks and wherein a first way has a fast access time; access control logic which manages access to the cache and is coupled to the plurality of ways; a plurality of multiplexers coupled to the first way in each of the banks and coupled to the access control logic; wherein the access control logic controls the multiplexer in a bank to remap any defective way in a bank to the first way in that same bank.

The inventive microprocessor die of the present invention comprises: self test logic which tests the die for defects; a set associative cache having a plurality of ways wherein the ways are segmented into a plurality of banks; access control logic which manages access to the cache coupled to the self test logic and coupled to the plurality of ways in said cache; a first way in the cache which has a physically shorter path to the access control logic; a plurality of multiplexers coupled to the first way in each of the plurality of banks and coupled to the access control logic; wherein the access control logic controls the multiplexer in a bank to remap any defective way in a bank to the first way in that same bank.

The method of absorbing defects in a set associative cache according to the present invention comprises: providing a set associative cache with a plurality of ways wherein the ways are segmented into a plurality of banks and wherein a first way has a fast access time; providing a plurality of multiplexers coupled to the first way in each of said banks; and using the multiplexer in a bank to remap any defective way in a bank to the first way in that same bank.

The computer system incorporating the present invention comprises: an output device to communicate information to a user; a microprocessor comprising: a set associative cache having a plurality of ways wherein the ways are segmented into a plurality of banks; access control logic which manages access to the cache coupled to the plurality of ways in said cache; a first way in the

cache which has a physically shorter path to the access control logic; a plurality of multiplexers coupled to the first way in each of the plurality of banks and coupled to the access control logic; wherein the access control logic can control the multiplexer in a bank to remap any defective way in a bank to the first way in that same bank.

5

BRIEF DESCRIPTION OF THE DRAWINGS

The invention can be more fully understood by referencing the accompanying drawings wherein:

Fig. 1 shows a block diagram of the architecture relating to a data array in set associative on-chip cache on a microprocessor die; and

Fig. 2 shows a block diagram of the architecture relating to a data array in set associative on-chip cache on a microprocessor die as contemplated by the present invention.

NOTATION AND NOMENCLATURE

15 Certain terms are used throughout the following description and claims to refer to particular system components. As one skilled in the art will appreciate, components may be referred to by different names. This document does not intend to distinguish between components that differ in name but not function. In the following discussion and in the claims, the terms “including” and “comprising” are used in an open-ended fashion, and thus should be interpreted to mean

20 “including, but not limited to...”. Also, the term “couple” or “couples” is intended to mean either an indirect or direct electrical connection. Thus, if a first device couples to a second device, that connection may be through a direct electrical connection, or through an indirect electrical

connection via other devices and connections. Finally, the term “logic” is used to encompass hardware and software solutions.

DETAILED DESCRIPTION OF THE DRAWINGS

5 Referring now to the drawings, wherein like reference characters denote corresponding components or parts:

Fig. 1 shows a functional block diagram of the architecture 10 relating to a data array of on-chip cache on a microprocessor die where the cache is configured as a 7-way set associative cache. The access control logic 12 is the portion of the microprocessor that controls, manages and performs the reads and writes to the cache data array 14. The data array 14 is in a standard set associative cache configuration with 7-ways and is segmented into four corners or banks 15. The data is written or stored across each bank 15 in one of the seven ways 16 as shown. Thus, when data is read from the data array of the cache, the data is read from a way 16 across all of the banks 15. Any number of banks 15 can be used with a plurality of ways 16 associated across the banks 15. The embodiment illustrated in Fig. 1 incorporates four banks 15 and seven ways 16 in each bank 15. The seven ways 16 in each bank 15 are designated numerically as way 0 through way 6 as shown. The four banks 15 are designated alphabetically as bank A-D. Although the embodiment shown comprises a set associative cache having seven ways 16 and four banks 15, it is understood that the data array 14 of the cache may be segmented with any granularity between banks 15 and ways 16. Typically, the data must be segmented in some format so that the data can be read out of the cache efficiently. Data stored in large monolithic data arrays takes longer to access and thus requires longer clock cycle times. The present invention applies to any cache formatted as a set associative cache regardless of granularity.

Each way 16 in each bank 15 is coupled to the access control logic 12 such that a set hit signal or signals 18 can be sent between the data array 14 and to the access control logic 12. More specifically, signals Hit 0 – Hit 6 are sent to way 0 – way 6 in each bank 15 of the data array 14, respectively. A set hit signal 18 is sent from the access control logic 12 to the specific way 16 in the cache data array 14 to which data is to be read or written.

Self-test logic 11 in the microprocessor is used to determine if there is any defective portion of the microprocessor die. In manufacturing, and then subsequently on each power-up of the microprocessor (*i.e.*, in a computer system when the power supply supplies power to the microprocessor), self-test logic built in to the processor tests for defects in the die, including in the data array of the cache. If a defect is found, the self-test logic 12 determines where the defect is located and takes appropriate corrective measures to repair the defect. Not all defects can be successfully repaired by the self-test logic 11. If the defect cannot be repaired, the location of all unrepairable defects is stored, typically in status registers. The location and number of unrepairable defects determines whether the die can be used or must be discarded.

The self-test logic 11 is coupled to the access control logic 12 both to perform the self-testing of the cache and to provide the results of the testing to the access control logic 12. As noted, generally the self test logic 11 stores the test results in status registers which the access control logic 12 can access to determine if there are any defective portions of the data array 14 of the cache. In a typical set associative cache, if there are any unrepairable defects in the data array 14, the entire way in which the defect is found must be disabled and unused. Otherwise, data stored in the defective way will be unreliable. Unfortunately, even if only one portion of the way (such as the portion of the way in one bank) were defective, normally the entire way would have to be disabled. Obviously, in a seven way associative cache such as the one shown, each defect in a

separate way would disable 1/7 of the effective and usable cache size. The present invention addresses this problem and provides an alternative method of working around defects in the data array while saving more of the overall cache size.

It should be noted that in any set associative cache configuration on a microprocessor die, one way 16 will be physically closer to the portion of the access control logic 12 having final control over the access to the data array 14 of the cache. In Fig. 1, way 0 is physically closer to the access control logic 12 while way 6 has the longest physical path to the access control logic 12. In any microprocessor configuration, there will be one way which has the shortest path. This difference in path lengths provides an opportunity.

The time required to access data in the cache is often the critical time for determining how fast the microprocessor can cycle. Thus, how fast the set hit signal 18 can be sent between the access control logic 12 and the data array 14 will often set the cycle time for the microprocessor. The cycle time must be set to accommodate the slowest way, that is, the way with the longest physical path, way 6 as illustrated in Fig. 1. Accordingly since the set hit signal 18 travels a shorter distance for way 0, there is additional time to perform additional functions in way 0 without affecting cycle time or without introducing any additional latency. Given this extra time available in the shortest path, a multiplexer or comparable switching mechanism (herein collectively referred to as a "mux") can be introduced into the path of the set hit signal 18 from way 0 (*i.e.*, the "Hit 0" signal) to the access control logic 12. So long as the additional time added by the mux (t_{mux}) does not make the access time for way 0 exceed the time for the longest way (t_{way6}) then the mux can be added without adding latency. This relationship can be expressed as:

$$t_{mux} + t_{way0} \leq t_{way6}$$

Alternatively, as long as the time added by the mux (t_{mux}) does not make the access time for way 0 exceed a clock cycle, then the mux can be added:

$$t_{mux} + t_{way0} \leq t_{clk}$$

Whether the time for the longest way (t_{way6}) or the clock cycle time (t_{clk}) is the critical parameter
5 depends on the system.

It should also be noted that a mux can be added in the path of any way, not just way 0, so long as the way has sufficient extra time to accommodate the added time of the mux (t_{mux}). Thus, a mux can be added to a path where the access time of the way (t_{wayn}) plus the time added by the mux (t_{mux}) does not exceed the time for the longest way (t_{way6}),

10 $t_{mux} + t_{wayn} \leq t_{way6}$

or alternatively, does not exceed a clock cycle:

$$t_{mux} + t_{wayn} \leq t_{clk}.$$

Incorporating muxes into multiple ways allows for even greater repair flexibility.

Fig. 2 shows a functional block diagram of the architecture 20 relating to the data array for
15 on-chip cache as contemplated by the present invention. Essentially, Fig. 2 illustrates the cache architecture of Fig. 1 with the addition of four multiplexers or muxes 22, 24, 26 and 28 in the path of way 0, that is, one mux in each path from way 0 in each of the four corners or banks 15. Note that a 7-way mux is used in each path because there are 7 ways in the embodiment shown. An n-way mux is required for an n-way set associative cache so that each set hit signal 18 for each
20 way can be muxed together in the path for way 0.

Fig. 2 also denotes unrepairable defects in certain ways with an "X" shown in the defective way. In particular, way 6 of bank A and way 5 of bank D are defective and are marked with an "X". In a traditional set associative cache, these defects would require disabling way 6 and way 5

in order to ensure data is not corrupted by storing it in these defective ways. Accordingly, 2/7 of the available cache size would be disabled due to the defects.

The addition of a mux 22, 24, 26, 28 in each of the way 0 paths allows for a more efficient work around of the defects shown. The work around is implemented as follows: Way 0 is
5 logically disabled so the access control logic 12 will not use way 0 for normal storage of data. With way 0 disabled for normal use, it can then be remapped and used as a surrogate for the defective ways in each bank using the muxes. In particular, mux 20 is set such that way 0 is used in place of defective way 6 in bank A. More specifically, the access control logic 12 sends a control signal to mux 20 such that a set hit signal 18 for way 6 (Hit 6 signal) is effectively
10 remapped to way 0 so that way 0 will be used in place of way 6 for bank A. Correspondingly, the portion of way 6 in bank A is disabled so that it will not attempt to put its data on the data bus at the same time as way 0 in bank A. This is accomplished by sending a disable signal to the portion of way 6 in bank A. The remaining portions of way 6, i.e., those portions in banks B, C, and D, remain active. So, data is effectively read from way 0 in bank A with way 6 in banks B, C, and D.

15 Similarly, way 0 in bank D can be remapped such that way 0 is used in place of defective way 5 of bank D by controlling mux 28 to remap the set hit signal for way 5 (Hit 5 signal) to way 0 for bank D and disabling way 5 of bank D. As a result of the muxes 22 and 28 then, when data is read from way 5, it is actually retrieved from way 5 in banks A, B and C and way 0 in bank D. Similarly, a read from way 6 actually retrieves data from way 6 of banks B, C and D and way 0 of
20 bank A. Thus, the portions of way 0 in the separate banks can be used in place of a defective way in each bank 15 without adding any latency to the system.

Note that in this example the defects have been successfully “absorbed” or “worked around” by only disabling one way, or 1/7 of the available cache size, instead of disabling 2 ways,

or 2/7 of the available cache, as would traditionally have been required. By extension, the muxes in the paths for way 0 can be used to remap around one defective way in each bank (with a maximum of four defective ways being remapped to way 0 in a four bank architecture as shown). Accordingly, if there were defects in way 2 of bank A, way 3 of bank B, way 4 of bank C, and way 5 of bank D, each defect could be remapped via muxes 22, 24, 26, 28 to way 0 in that bank. Thus, instead of having to disable four separate ways to work around the defects, 4/7 of the cache, all of the defects can be absorbed with only way 0 disabled for normal use, 1/7 of the available cache.

It is understood that there are limitations to the embodiment as described. For instance, if there are two defective ways in one bank, then only one of the defective ways can be remapped to way 0. In addition, if a defect occurs in way 0 then no defects in that bank can be remapped to way 0. Finally, if there is only one defective way, remapping that way to way 0 results in no savings since one way would still have to be disabled, *i.e.*, 1/7 of the cache. Some of these limitations can be overcome, however, by placing muxes in the second (or more) shortest way, assuming its access time is fast enough such that there is enough extra time to accommodate the added time of the mux as discussed above. Having two or more ways with muxes incorporated in their paths would allow multiple defective ways in the same bank to be remapped.

Ultimately, depending on the number and location of the defects in the data array, the invention allows for more defects to be absorbed in the data array without sacrificing as much of the total available cache. This can result in resurrecting useful parts that would have been discarded in the past, thereby increasing the overall manufacturing yield for the microprocessor die. Although the invention may be used to work around more defects in the data array of the microprocessor and thus increase manufacturing yields by allowing more useable parts to be shipped, the invention is also useful for debug of the microprocessor. In particular, the invention

can allow earlier debug of the microprocessor because you do not have to wait until the manufacturer has debugged the manufacturing process to obtain parts having approximately a full on-chip cache available for testing. This debug advantage alone may warrant the addition of the invention to the architecture of a microprocessor.

- 5 The above discussion is meant to be illustrative of the principles and various embodiments of the present invention. While the invention has been particularly shown and described with respect to specific embodiments thereof, numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.